

Multi-Agent based simulation of FOREX exchange market

Vivien Delage^{a b}, Christian Brandlhuber^b, Karl Tuyls^a, and Gerhard Weiss^a

^a *Department of Knowledge Engineering, Maastricht University, the Netherlands*
^b *Teramark Technologies GmbH, Munich, Germany*

Abstract

The FOREX market is quite unique among the traditional financial markets. First, with a very high transaction volume, it is the most liquid market of the world. Second, this market is influenced by a vast and diverse panel of information. This ranges from macro-economic elements such as interest rate parity, inflation rate or unemployment rate to political elements such as changes in government. Third, the decentralized structure of this market makes its study and interpretation very challenging. This paper presents a distributed multi-agent based model for simulating the behaviour of the FOREX market. The model allows the replication of the market price and the study of the complex structure of the FOREX market. The multi-agent model is described and validated by extensive experiments.

1 Introduction

Financial markets consist of a collection of buyers and sellers of one or more goods, as well as the transactions between them. In such a market, traders can exchange stock, bonds, commodities, and other items at low transaction costs. Markets are typically governed by rules specifying what traders can do (e.g., which messages they can exchange) and how commodities are allocated given their actions. The most familiar markets are centralized ones, such as auctions [14], in which messages from traders, e.g., offers to buy or sell at a given price, are sent to a central unit that processes the trades. However, markets can also be decentralized, in which case traders directly buy and sell from each other.

In recent years, the structures of the financial markets that comprise the global economy have changed dramatically [9-13]. Rather than individual, segregated markets using centralized auctions (e.g., as in a traditional stock exchange), new partially decentralized markets have emerged. These markets are often built on individual trading relationships between financial institutions and thus transcend industries and borders. Together, these relationships form ad-hoc networks of unprecedented size and complexity. An example of such a market is the foreign currency exchange market or FOREX, the type of market under investigation in this article.

While such networked markets create new opportunities and efficiencies, they also lack a principled manner of modelling them. In fact, the recent global economic turmoil has demonstrated how events in one market, such as the U.S. sub-prime mortgage market, can have profound effects across the world. Consequently, an increasingly urgent question is how to model such markets and how to develop algorithms for automated trading in this type of distributed market. Unfortunately, existing models, e.g., those based on dynamic stochastic general equilibria [15] are poorly suited to modelling such distributed markets. Furthermore, many algorithmic trading strategies have been proposed for modelling traders, such as Gjerstad Dickhaut [16], Roth-Erev [17] and Zero Intelligence Plus [6], which are not applicable in inter-bank markets because they assume access to historical transaction information. Since inter-bank markets are partially decentralized, such transaction information is hidden from the trading agents, who can view only the price feed. In this paper we present a distributed agent model of FOREX and propose a modified version of Zero Intelligence Plus that bases its actions only on such price information. The new

model is specifically designed with the market decentralization in mind. It uses a layered architecture inspired by microstructure approach [2] in order to map at best the structure of the FOREX market.

The remainder of the paper is structured as follows. In Section 2 the necessary background on FOREX and trading strategies is provided in order to understand the remainder of the article. Section 3 presents the new model architecture as well as the different type of market agents in use. Section 4 provides results of the experiments conducted to validate the new model. Finally section 5 concludes the paper with a discussion of possible model extensions.

2 Background

The foreign currency exchange market (FOREX) differs substantially from traditional markets such as stock exchanges in multiple aspects: First, it is a partially decentralized market: there is no central trading authority. Instead the market is layered into an inter-bank market and a retail market. The inter-bank market is where transactions take place between major banks and big institutions. The retail market is where regular/small traders deal through the banks (the banks being the dealers). Second, it has a continuous operation: the market is running 24h a day with no interruption except during weekends. Third, the use of high leverage is very common. For example, a trader could have 10M \$ on account and be able to trade tickets of the size of 100M \$. Fourth, it has a huge trading volume leading to very high liquidity (estimated around 4 trillion \$ average daily turnover [1]).

Over the past years the FOREX market has changed its architecture from a fully decentralized market (the banks were dealing with each other directly through the phone) to a partially centralized market where most of the inter-bank transactions are done through central electronic brokers (e.g. EBS or Reuters). However, the retail traders still do not have access to the central exchange place, only the banks and very big institutions have access to it. This is due to the fact that connecting to this central exchange place is very expensive and that the minimum deal size on this market is very high, thus it is not suitable for small players.

Because of the decentralization of the market, all participants do not share the same information. Indeed the banks have access to private information in the form of the order flows coming from their customers [2]. Since the banks are taking both the role of a dealer and the role of a trader in the market, this gives them a major advantage when issuing speculative trades.

In accordance with the architecture, the market shares are also in constant evolution since the past few years. Deals initiated in the inter-bank market have gone from more than 60% of the global market turnover in 1995 to less than 50% in 2010 [1]. This is not due to a decrease in the inter-bank market activity but rather to a constant increase of deals made by other institutions such as hedge funds, pension funds or insurance companies.

2.1 The macro-economic model

Macro-economists have been trying to model the FOREX market using different approaches, the most common one being the asset approach which is driven by a set of macro fundamentals. In this case the market is modelled using the following equation:

$$\Delta P_t = f(i, m, z) + \varepsilon_t$$

Where ΔP_t is the change in the exchange rate over the time period, i is the interest rate parity between the exchanged currencies, m is the money supply, z representing other macro determinants (e.g. inflation rate, unemployment data, etc) and ε_t representing the residual noise term.

These fundamentals are known to have an influence on the FOREX market. Indeed each currency is tightly linked to its country macro-economic and political states. However, what is interesting to note is that none of the macro-economic models developed so far succeeds to explain FX rates levels and volatility in the long term. The microstructure approach [2] attributes this failure to three main reasons:

- Macro-economic models forget to take into consideration that not all information relevant for trading is publicly available (private information exists).

- Different market players have different effects on the price (funds do not have the same impact on price as private traders or central banks). The macro-economic models only consider one global price effect.
- Different trading mechanisms have different effects on the price (e.g. inter-bank dealing mechanism is different to dealer-customer mechanism of the retail market). This is also not taken into consideration by macro-economic models.

2.2 Zero intelligence based models

With his experiment conducted in 1955, Smith [3] was already showing that it is possible for a market to converge to equilibrium even with a modest number of inexperienced traders having no prior knowledge of the market. So he proved that by using a double auction mechanism [4], even traders that are not rational players (in the sense that they do not have knowledge of supply and demand) will still lead the market price to equilibration.

Gode and Sunder (1993) conducted a similar experiment [5] but this time replacing the human traders by automated trade agents. Each of the automated traders was given a private value and was instructed to submit random bids and asks constrained to the private value they had. The goal of this experiment was to check whether automated agents without any type of rationality (Zero Intelligence) could lead the market to converge toward the same equilibrium as the human traders. Gode and Sunder's finding is that indeed the automated agent could replicate the same kind market behaviour as the human traders. They explain this result by the discipline imposed by the use of a continuous double auction (CDA) process.

In his work from 1997, Dave Cliff criticized the work of Gode and Sunder. He acknowledged that the structure of the double auction process is responsible for the high market efficiency without the need for agent rationality. He questioned, however, the role of the CDA in the tendency of the market to converge to an equilibrium price. He was then demonstrating that the convergence to equilibrium is only happening as a result of the specific market setup Gode and Sunder were using and that it does not hold in general. He then introduced a new version of the automated trade agent called Zero Intelligence Plus (ZIP). The ZIP agent is also given a private value for trading. It uses a basic learning algorithm in order to adapt its profit expectation by monitoring the actions of the other trade agents of the market. These changes make the ZIP agent more complex but it ensures that the simulated market price will converge toward equilibrium in all cases. Further, since the agent has a memory the equilibrium price will stay stable over the different trading periods (see. figure 1) and there will be no reset at the beginning of each new trading period as it was the case with the original ZI agent.




Figure 1: Typical transaction graph using ZIP agents over 6 trading periods

In the mean time, improved versions of the ZIP agent have been published such as the ZIP60 agent [7] and the GD agents [16]. Those new agents focus on improving the profitability, but they also use more complex learning algorithms.

All of those automated trade agents have good performance; however they have been designed to be used in a centralized market. Indeed, they are assumed to be connected with a central double auction market where all other market players are visible. They also suppose that all transactions and price shouts are publically available through the CDA process. These assumptions do not hold for the FOREX market.

As we have already argued FOREX is partially decentralized. Due to this the usual traders do not have access to transaction information. The only information they receive is the price feed coming from the dealer (the banks) they are connected to. In addition, because of decentralization, the market participants are scattered and are not visible to each other.

3 Multi-agent decentralized model

Trying to resolve the major flaws of macro-economic models, a new model is proposed in this paper by combining techniques from the microstructure approach with multi-agent technology. The new model as presented in Figure 2 attempts to map the specificities of the FOREX market. It is composed of three different types of agents: a CDA process, an agent representing the bank and a trader agent.

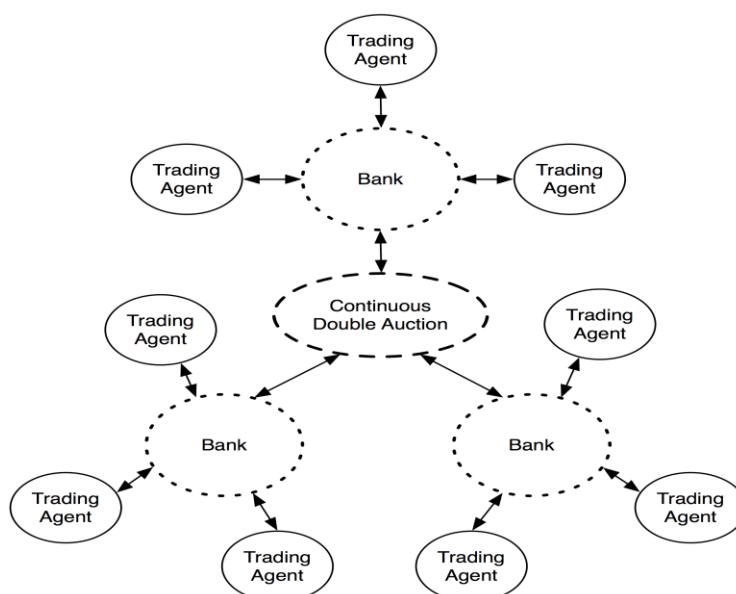


Figure 2: Model of the partially decentralized market.

The CDA in the centre of the model is representing the inter-bank market. This role is usually taken by electronic brokers such as EBS or Reuters in the real market. The CDA is responsible for the match-making between the bank agents. When a bid and an offer cross each other the CDA creates a transaction whose price is the average between both shouts and informs all connected agents of the transaction. In this case this side of the market is transparent as each connected agent will be informed of all shouts and transactions taking place on the CDA. The CDA also keeps an order book listing all shouts that have not yet been matched. This order book is queried when a new shout arrives to determine whether a deal can be made, if not the new shout is added to the order book. The order book has the same life time as one trading period and will be flushed when the trading period ends.

The bank agent has a double role. It first acts as a dealer for his customers (the trader agents). In this role the match-making process is quite similar to the way the CDA is working. However there is a notable difference: when a bid and an offer cross each other the bank creates a transaction matching exactly the price given by the trade agents. It then takes the spread between the two shouts as a profit for itself. The second role a bank takes is as a trader. Indeed, when dealing with his customers the bank is building up a position inventory. Bids and asks coming from the trade agents will be first matched internally by the bank. However, in case of a surplus on one side (long or short) the bank will need to clear its inventory in order to reduce the trading risk. It will do so by sending any excess in its inventory to the CDA process to be exchanged with other banks. The bank can also exercise its trader role as a pure speculator, trying then to use the private information it got from its customer base as its own advantage.

The trader agent used in this model is a modified version of the original ZIP agent. As we have seen, the ZIP agent was not designed to be used in a decentralized market. As a result the learning algorithm in place is assumed to have access to the entire set of transaction information occurring in the market. In this case the transaction information is not available to the trade agent thus the learning algorithm has to be tweaked so that it only uses the price information provided by the bank as an input. Just as with the ZIP agent, the modified trader is given a private value at start-up in addition to a side for his trade (a trade agent can either be a buyer or a seller agent but cannot act as both). Below is the pseudo-code of the modified version of ZIP agent algorithm:

For each price quote q (with bid component q_b and ask component q_a) coming from the bank:

- Any sellers S_i whose current profit margin $P_i \leq q_b$ should raise their margin
- Any active sellers S_i whose current profit margin $P_i > q_b$ should lower their margin
- Any buyers B_i whose current profit margin $P_i \geq q_a$ should raise their margin
- Any active buyers B_i whose current profit margin $P_i < q_a$ should lower their margin

The model is composed of one CDA process and a fixed set of bank agents. This is not the case with the trader agents (both buyers and sellers) whose number will vary during the simulation. The idea is to adjust at runtime the ratio between the number of buyer and seller agents in order to control the price produced by the simulation. The generated price needs to be structured in a way that it replicates as closely as possible the real market price.

A good estimate of the buyer/seller ratio is needed if we want the model to have good performance with market replication. To do so a Bayesian network is set up as displayed in figure 3. The Bayesian network [8] is trained using samples coming from real EUR/USD price data gathered over the period from 2005 to 2010. It is then queried during the simulation to provide an estimate for the ratio.




Figure 3: Bayesian network as buyer/seller ratio estimator.

The inputs to the Bayesian network are mapping similar fundamentals used by macro-economists. They contain the interest rate parity between the two countries owning the asset currencies. They also contain the inflation parity between the same two countries. And finally they contain data coming from the commitment of traders (COT) reports [9].

The COT reports are publicly available information published by the US commodity future trading commission at the end of each trading week. These reports give information about positions being held in variety of traded goods on the different US exchange places. This includes major currency indexes such as USD and Euro indexes. The data gives a snapshot of the overall long and short positions being held by 2 different trading groups (commercial and non-commercial). It is intended to provide a good estimate on volume traded during the week. This information is otherwise hidden in the market, and thus it provides valuable inside on order flow information.

4 Experiments

In order to validate the new model a set of experiments are conducted. Each experiment runs as a sequence of distinct trading periods or “rounds”. In the beginning of each round all previous price shouts that have not been matched are cleared and the CDA restarts from scratch. To simulate a real time market, the CDA/banks use a random round-robin algorithm to randomly pick a trader. A picked agent is then informed of the latest prices and asked if it wants to shout a new order in the market or not. Once all trade agents have been picked the picking algorithm restarts from the beginning. This continues until no more transactions are happening between agents. If this is the case the round is considered as finished and a new trading period can start.

4.1 Market equilibration

The first experiment is set up to check whether the new decentralized model together with the modified version of the ZIP agent is still able to lead the market price to an equilibrium state. The following architecture is applied: 5 bank agents, and each bank agent is connected to a set of 200 ZIP traders (100 buyers and 100 sellers). The number of trade agents remains static during the simulation. The simulation running over 100 trading periods produces the following transaction graph displayed in figure 4.




Figure 4: Generated transaction price on the bank level.

A study of the graph shows that an equilibrium price is indeed found during the simulation (around a price of 100). However comparing these results with those obtained by Dave Cliff with the original ZIP agent (cf. Figure 1), we can observe that the overall price volatility is much higher in this experiment. Even after reaching the equilibrium the price is still oscillating around it for a very long time. Further, the initial learning phase takes much longer before the price stabilizes compared to the original test. This is explained by the fact that the amount and quality of information given to the new ZIP agent is much less than originally. So the ZIP agents acting with incomplete information will decrease the market stability and increase the overall volatility.

4.2 Market replication

The second experiment is set up in order to test the replication capability of the new model versus real market data. The asset chosen is EUR/USD over the period from year 2005 to 2010. A trading period is set to be equivalent to 1 week in real time. The number of bank agents is kept to 5 but this time the number of buyer and seller agents vary depending to the ratio given by the Bayesian network. The simulation runs five consecutive times and produces the following transaction graph displayed in figure 5. For comparison the real price data is also added to the graph.

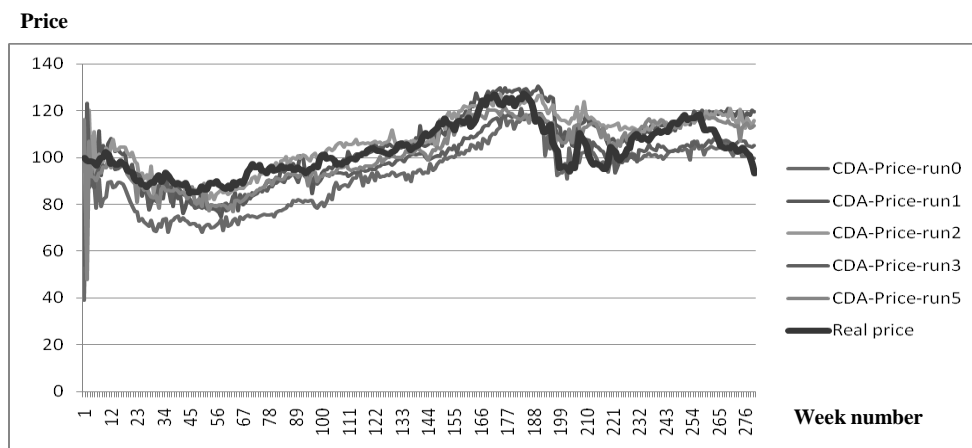


Figure 5: Comparison between generated transaction prices and real market price.

By inspecting the graph we can say that in all cases the simulation is able to replicate the real market price with a comparably good quality. If we omit the very early period, where the ZIP agents are still in the learning phase, the generated curves follow the same trend as the real price quite closely.

5 Conclusion

The FOREX market is a challenging market due to its decentralized architecture which makes it more complex to interpret than typical centralized markets. We presented a new multi-agent based model intended to simulate the behaviour of FOREX. The model consists of two layers in order to replicate both the inter-dealer and the retail components of the foreign currency exchange market. We also introduced a modified version of the ZIP trade agent. As transaction information is not available in the decentralized market, the modified trade agent bases its actions solely on price information available from the banks. The new model was then tested and was able to lead the simulated market into an equilibrium state. The model was also analyzed through a replication test against the EUR/USD asset pair and was able to replicate the FOREX market with a good performance over a period of 5 years. The replication was achieved by controlling the ratio between the buyer and seller agents using a Bayesian network. The model results are very promising and show a good potential for improvements:

On the trade agent side, in the current model only one type of agent is used, namely, the modified ZIP agent. It would be interesting to integrate other types of trade agents such as the ZIP60 agent or the GD agent and to analyze how they affect the simulated market behaviour. Furthermore the model would profit from adding different classes of trade agents. We know that the real market contains different types of traders such as hedgers, speculators or the central banks. Therefore it would be useful to have trade agents in the model that directly map such roles and to explore how these agents interact with each other.

Concerning the bank agent, the dealer algorithm it uses in the current model is quite basic. Indeed it only uses basic risk-averse rules in order to liquidate its open position as soon as possible. An important extension would be to implement more complex rules and strategies which not only take into consideration the order flow coming from its customer but also the actions of the other bank agents. It

would also be important to make the bank agent able to enter into speculative trades. This feature is currently missing in the model and result in a higher fidelity compared to the real market.

Finally, we have seen that the current model was able to replicate the real market. A natural next step would be to use the multi-agent model to try to predict the market price. This could be achieved by training the Bayesian network on a preliminary set of data and applying it on a set of unknown data.

References

- [1] Triennial Central Bank Survey from the Bank for international statement, 2010: <http://www.bis.org/publ/rpx10.pdf>
- [2] Richard K. Lyons: The microstructure approach to exchange rates. MIT Press, 2001
- [3] Smith: An Experimental Study of Competitive Market Behavior. University of Chicago Press, 1962
- [4] Simon Parson: Everything you wanted to know about double auctions, but were afraid to (bid or) ask, 2006
- [5] Dhananjay K. Gode and Shyam Sunder: Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. University of Chicago Press, 1993
- [6] Dave Cliff: Minimal-Intelligence Agents for Bargaining Behaviors in Market-Base Environments. HP Labs Technical Reports, 1997
- [7] Dave Cliff: ZIP60: Further explorations in the Evolutionary Design of Online Auction market Mechanisms. IEEE Transactions on Evolutionary Computation, 2005
- [8] Niculesc, Mitchell, Rao: Bayesian Network Learning with Parameter Constraint. The Journal of Machine Learning Research Volume 7, 2006
- [9] CFTC – commitment of traders reports (every end of trading weeks): <http://www.cftc.gov/marketreports/commitmentsofraders/index.htm>
- [10] R. A. Cagliano, M. D. Fraser, and Mark E. Shaefer. Auction allocation of computing resources. Communications of the ACM, 38:88–102, 1995.
- [11] New York Stock Exchange. TStock market activity report. http://www.nyse.com/pdfs/02_STOCKMARKETACTIVITY.pdf, 2002.
- [12] P. Cramton, A. E. Khan, R. H. Porter, and R. D. Tabors. Pricing in the california power exchange electricity market: Should california switch from uniform pricing to pay-as-bid pricing? Technical report, California Power Exchange, California, USA, 2001.
- [13] O. Kephart. Software agents and the route to the information economy. Proceedings of the National Academy of Sciences, 99:7207–7213, 2002.
- [14] D. Friedman and J. Rust. The Double Auction Market: Institutions, Theories and Evidence. Santa Fe Institute Studies in the Sciences of Complexity. Perseus Publishing, 1993.
- [15] S. Parsons, J. A. Rodriguez-Aguilar, and M. Klein. Auctions and bidding: A guide for computer scientists. ACM Computing Surveys, to appear, 2011.
- [16] Steven Gjerstad and John Dickhaut. Price formation in double auctions. Games and Economic Behaviour, 22(1):1–29, 1998.
- [17] I. Erev and A.E. Roth. Predicting how people play games with unique, mixed-strategy equilibria. American Economic Review, 88:848–881, 1998.